


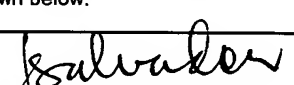
10-06-05


EPW  
#

PTO/SB/21 (09-04)

<b>PTO TRANSMITTAL FORM</b> OCT 04 2005 (to be used for correspondence after initial filing) PATENT & TRADEMARK OFFICE	Application Number	10/719,286
	Filing Date	November 20, 2003
	First Named Inventor	Fujibayashi, Akira
	Art Unit	3184
	Examiner Name	Unassigned
	Attorney Docket Number	16869B-077700US
Number of Pages in This Submission		17

ENCLOSURES (Check all that apply)		
<input checked="" type="checkbox"/> Petition Fee Transmittal <input type="checkbox"/> Fee Attached <input type="checkbox"/> Amendment/Reply <input type="checkbox"/> After Final <input type="checkbox"/> Affidavits/declaration(s) <input type="checkbox"/> Extension of Time Request <input type="checkbox"/> Express Abandonment Request <input type="checkbox"/> Information Disclosure Statement  <input type="checkbox"/> Certified Copy of Priority Document(s) <input type="checkbox"/> Reply to Missing Parts/ Incomplete Application <input type="checkbox"/> Reply to Missing Parts under 37 CFR 1.52 or 1.53	<input type="checkbox"/> Drawing(s) <input type="checkbox"/> Licensing-related Papers <input checked="" type="checkbox"/> Petition to Make Special <input type="checkbox"/> Petition to Convert to a Provisional Application <input type="checkbox"/> Power of Attorney, Revocation <input type="checkbox"/> Change of Correspondence Address <input type="checkbox"/> Terminal Disclaimer <input type="checkbox"/> Request for Refund <input type="checkbox"/> CD, Number of CD(s) _____ <input type="checkbox"/> Landscape Table on CD	<input type="checkbox"/> After Allowance Communication to TC <input type="checkbox"/> Appeal Communication to Board of Appeals and Interferences <input type="checkbox"/> Appeal Communication to TC (Appeal Notice, Brief, Reply Brief) <input type="checkbox"/> Proprietary Information <input type="checkbox"/> Status Letter <input checked="" type="checkbox"/> Other Enclosure(s) (please identify below): Return Postcard Ten (10) cited references
Remarks: The Commissioner is authorized to charge any additional fees to Deposit Account 20-1430.		
SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT		
Firm Name	Townsend and Townsend and Crew LLP	
Signature		
Printed name	Chun-Pok Leung	
Date	October 4, 2005	Reg. No. 41,405

CERTIFICATE OF TRANSMISSION/MAILING			
Express Mail Label: EV 529871546 US			
I hereby certify that this correspondence is being deposited with the United States Postal Service with "Express Mail Post Office to Address" service under 37 CFR 1.10 on this date <b>October 4, 2005</b> and is addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below.			
Signature			
Typed or printed name	Joy Salvador	Date	October 4, 2005



**PETITION FEE**  
**Under 37 CFR 1.17(f), (g) & (h)**  
**TRANSMITTAL**  
(Fees are subject to annual revision)

Send completed form to: Commissioner for Patents  
P.O. Box 1450, Alexandria, VA 22313-1450

Application Number	10/719,286
Filing Date	November 20, 2003
First Named Inventor	Fujibayashi, Akira
Art Unit	3184
Examiner Name	Unassigned
Attorney Docket Number	16869B-077700US

Enclosed is a petition filed under 37 CFR §1.102(d) that requires a processing fee (37 CFR 1.17(f), (g), or (h)). Payment of \$ 130.00 is enclosed.

This form should be included with the above-mentioned petition and faxed or mailed to the Office using the appropriate Mail Stop (e.g., Mail Stop Petition), if applicable. For transmittal of processing fees under 37 CFR 1.17(i), see or PTO/SB/17i.

**Payment of Fees** (small entity amounts are NOT available for the petition fees)

- ☒ The Commissioner is hereby authorized to charge the following fees to Deposit Account No. 20-1430 :  
☒ petition fee under 37 CFR 1.17(f), (g) or (h) ☒ any deficiency of fees and credit of any overpayments  
Enclose a duplicative copy of this form for fee processing.

☐ Check in the amount of \$ \_\_\_\_\_ is enclosed.

☐ Payment by credit card (Form PTO-2038 or equivalent enclosed). Do not provide credit card information on this form.

**Petition Fees under 37 CFR 1.17(f): Fee \$400 Fee Code 1462**

For petitions filed under:

- § 1.53(e) - to accord a filing date.
- § 1.57(a) - to accord a filing date.
- § 1.182 - for decision on a question not specifically provided for.
- § 1.183 - to suspend the rules.
- § 1.378(e) - for reconsideration of decision on petition refusing to accept delayed payment of maintenance fee in an expired patent.
- § 1.741(b) - to accord a filing date to an application under § 1.740 for extension of a patent term.

**Petition Fees under 37 CFR 1.17(g): Fee \$200 Fee Code 1463**


For petitions filed under:

- § 1.12 - for access to an assignment record.
- § 1.14 - for access to an application.
- § 1.47 - for filing by other than all the inventors or a person not the inventor.
- § 1.59 - for expungement of information.
- § 1.103(a) - to suspend action in an application.
- § 1.136(b) - for review of a request for extension of time when the provisions of section 1.136(a) are not available.
- § 1.295 - for review of refusal to publish a statutory invention registration.
- § 1.296 - to withdraw a request for publication of a statutory invention registration filed on or after the date the notice of intent to publish issued.
- § 1.377 - for review of decision refusing to accept and record payment of a maintenance fee filed prior to expiration of a patent.
- § 1.550(c) - for patent owner requests for extension of time in ex parte reexamination proceedings.
- § 1.956 - for patent owner requests for extension of time in inter partes reexamination proceedings.
- § 5.12 - for expedited handling of a foreign filing license.
- § 5.15 - for changing the scope of a license.
- § 5.25 - for retroactive license.

**Petition Fees under 37 CFR 1.17(h): Fee \$130 Fee Code 1464**

For petitions filed under:

- § 1.19(g) - to request documents in a form other than that provided in this part.
- § 1.84 - for accepting color drawings or photographs.
- § 1.91 - for entry of a model or exhibit.
- ☒ § 1.102(d) - to make an application special.
- § 1.138(c) - to expressly abandon an application to avoid publication.
- § 1.313 - to withdraw an application from issue.
- § 1.314 - to defer issuance of a patent.

  
\_\_\_\_\_  
Signature  
**Chun-Pok Leung**  
\_\_\_\_\_  
Typed or printed name

\_\_\_\_\_  
October 4, 2005  
Date  
**41,405**  
\_\_\_\_\_  
Registration No., if applicable



PATENT  
Attorney Docket No.: 16869B-077700US  
Client Ref. No.: HAL267  
(340300827US01)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

AKIRA FUJIBAYASHI

Application No.: 10/719,286

Filed: November 20, 2003

For: METHOD AND APPARATUS  
FOR VOLUME REPLICATION  
MANAGEMENT AT PLANNED  
AND UNPLANNED LINK  
DOWN

Customer No.: 20350

Examiner: Unassigned

Technology Center/Art Unit: 2171

Confirmation No.: 3184

**PETITION TO MAKE SPECIAL FOR  
NEW APPLICATION UNDER M.P.E.P.  
§ 708.02, VIII & 37 C.F.R. § 1.102(d)**

Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

This is a petition to make special the above-identified application under MPEP § 708.02, VIII & 37 C.F.R. § 1.102(d). The application has not received any examination by an Examiner.

(a) The Commissioner is authorized to charge the petition fee of \$130 under 37 C.F.R. § 1.17(i) and any other fees associated with this paper to Deposit Account 20-1430.

(b) All the claims are believed to be directed to a single invention. If the Office determines that all the claims presented are not obviously directed to a single invention, then Applicants will make an election without traverse as a prerequisite to the grant of special status.

10/07/2005 HLE333 00000019 201430 10719286

01 FC:1464 130.00 DA

(c) Pre-examination searches were made of U.S. issued patents, including a classification search and a key word search. The classification search was conducted on or around June 30, 2005 covering Class 707 (subclass 204), Class 709 (subclasses 223 and 226), Class 710 (subclasses 22 and 23), Class 711 (subclasses 112, 114, 141, 161, and 162), and Class 714 (subclasses 5-7 and 15), by a professional search firm, Mattingly, Stanger, Malur & Brundidge, P.C. The key word search was performed on the USPTO full-text database including published U.S. patent applications.

(d) The following references, copies of which are attached herewith, are deemed most closely related to the subject matter encompassed by the claims:

- (1) U.S. Patent No. 5,504,861;
- (2) U.S. Patent No. 5,835,953;
- (3) U.S. Patent No. 5,857,208;
- (4) U.S. Patent No. 5,928,367;
- (5) U.S. Patent No. 6,260,124 B1;
- (6) U.S. Patent No. 6,543,001 B2;
- (7) U.S. Patent No. 6,578,120 B1;
- (8) U.S. Patent Publication No. 2003/0126107 A1;
- (9) U.S. Patent Publication No. 2003/0135650 A1; and
- (10) U.S. Patent Publication No. 2003/0177321 A1.

(e) Set forth below is a detailed discussion of references which points out with particularity how the claimed subject matter is distinguishable over the references.

A. Claimed Embodiments of the Present Invention

The claimed embodiments relate to systems for storing data, and in particular to volume replication in storage area networks. A technique is provided for restoring data after suspension of a communications link between two storage systems. Upon suspension of the link, an image of the data at the first and second locations is created in a secure location.

While the link is down, updates to the data at each of the first and second locations are maintained separately. Upon reestablishment of a communications link, the two locations are again synchronized.

Independent claim 1 recites, in a system having a primary storage volume having first data stored thereon and a secondary storage volume having second data stored thereon, the primary storage volume coupled to the secondary storage volume by a communications link, a method of restoring data after a suspension of the link. The method comprises upon suspension of the link, maintaining a first image of the first data stored on the primary volume and maintaining a second image of second data stored on the secondary volume; tracking updates to the first data and updates to the second data after suspension of the link; revising the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored; revising the image of the second data to account for the updates to thereby provide a complete second data image for use when the link is restored; and copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume.

Independent claim 7 recites a system for restoring data after an interruption in a communications link. The system comprises a primary storage volume having first data stored thereon; a secondary storage volume having second data stored thereon, the primary storage volume being coupled to the secondary storage volume by the communications link; a first storage controller for maintaining a first image of the first data stored on the primary volume upon the interruption in the communications link; a second storage controller for maintaining a second image of the second data stored on the secondary volume upon the interruption in the communications link; a first update storage for storing updates to the first data after the interruption in the communications link; a second update storage for storing updates to the second data after the interruption in the communications link; the first storage controller revising the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored; the second storage controller revising the image of the second data to account for the updates to thereby provide a complete second data image for use when the link is restored; and at least one of the first and

the second storage controllers revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored.

Independent claim 12 recites a sub-system for restoring data after an interruption in a communications link. The sub-system comprises a primary storage volume having first data stored thereon; a first storage controller for maintaining a first image of the first data stored on the primary volume upon the interruption in the communications link; a first update storage for storing updates to the first data after the interruption in the communications link; and the first storage controller revising the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored.

One of the benefits that may be derived is a robust technique for restoring data after suspension of a communications link between two storage systems.

B. Discussion of the References

1. U.S. Patent No. 5,504,861

The patent to Crockett et al., US 5504861, discloses a remote data shadowing system that provides storage based, real time disaster recovery capability. Record updates at a primary site 14 cause write I/O operations in a storage subsystem therein. The self describing record sets are transmitted to a remote secondary site 15 wherein consistency groups are formed such that the record updates are ordered so that the record updates can be shadowed in an order consistent with the order the record updates cause write I/O operations at the primary site. See column 12, lines 47-50. The record updates are written according to full consistency group recovery rules such that should the primary site be unavailable, the secondary site can recover a consistency group. See column 17, lines 31-47. The secondary site forms a "duplex pair" with the primary site. The duplex pair can become "failed" due to communication failure between the primary storage controller 3 and the secondary storage controller 6 via communication links 8. Alternatively, duplex pair can become "failed" due to errors in either the primary or secondary subsystem. If the failure is in the communication links 8, then the primary storage controller 3 is unable to communicate the failure directly to the secondary storage controller 6. The I/O error recovery program 2 quiesces the application programs hence taking control of the primary processor 1 for error recovery and data integrity

before returning control to the application requesting the write I/O operation. See column 7, line 57 to column 8, line 6.

Crockett et al. is directed to remote data duplexing. Although it discloses record updates and consistency in recovery, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Crockett et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

2. U.S. Patent No. 5,835,953

The patent to Ohran, US 5835953, discloses a system and method for maintaining logically consistent backups using minimal data transfer. A system comprises a backup system 14 having a backup storage device and one or more primary systems 12 having mass storage devices 20 that are to be backed up on the backup storage device. The primary systems identify changes that are going to be made to the mass storage device. The combined effected locations in the mass storage device of these identified changes are then captured in a static snapshot when the mass storage device is in a logically consistent state. See column 10, line 65 to column 11, line 5. Only those data blocks changed since the last backup are then transferred to backup system. The backup system can then store these changes or apply the changes to the backup storage device in order to bring the backup storage device current to a particular point in time. If a single record of the database was changed ten times between the time the last backup was made and the current backup time, certain prior art systems would send ten changes to the backup storage device. Ohran, however, simply sends the last change that was made before the current backup time. This

reduces the amount of data sent to the backup device to the very minimum needed to make a logically consistent backup. See column 12, lines 11-34.

Ohran is directed to maintaining logically consistent backups using minimal data transfer. Although it discloses updating changes in a logically consistent state, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Ohran fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

3. U.S. Patent No. 5,857,208

The patent to Ofek, US 5857208, discloses a data network with a remote data facility for providing redundant data storage and for enabling concurrent point-in-time backup operations. A local data processing system 10 with a data facility stores a data base and processes applications. A second system 11, physically separated from the first system, includes a data facility that normally mirrors the data in the first system. In a backup mode, the second system is enabled to transfer data from its data facility to a backup facility concurrently with, but independently of, the operation of the first system. See column 8, lines 7-58. On completion of the backup operation, the second system reconnects with and synchronizes with the first system thereby to reestablish the mirroring operation of the second system. See column 8, line 64 to column 9, line 40. When it is necessary to obtain a backup, that operation occurs at the remote system 11 concurrently with the continued operations within the local system 10 and without any intervention by the local system 10 that could adversely affect its operating characteristics. Immediately upon completion of the backup, the local and remote systems resynchronize to reestablish a mirror relationship. Typically the



number of tracks that need to be updated will be minimal, so that the time required to resynchronize the system after running decision support system applications will be minimal. See column 10, line 60 to column 11, line 10.

Ofek is directed to a point in time backup operation. Although it discloses an independent backup transfer of data, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Ofek fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

4. U.S. Patent No. 5,928,367

The patent to Nelson et al., US 5928367, discloses a disk storage control system 10 includes dual controllers 20, 25 having real-time, synchronous, mirrored memory therebetween to provide immediate, accurate, and reliable failover in the event of controller failure. Non-volatile random access memory 30, 35 provides retention of data during a loss of power and during the manipulation of hardware for purposes of repair. A communication path is established within the mirrored memory 30, 35 between the controllers to monitor and coordinate their activities. See column 2, line 48 to column 3, line 5. The state of the mirrored memory is continuously monitored for accuracy of the mirror and failure detection. Data consistency of the mirror image on each controller memory 30, 35 is ensured by "scrubbing" of the memory. See column 11, lines 32-37. In the dual controller context, the logical Master controller performs the scrub process preferably in a background mode or when the system 10 is idle. The scrubbing process reads each and every address location to find correctable/uncorrectable errors and to fix such errors to ensure consistency between the controllers. Scrubbing continually monitors the consistency of the two mirrored memory

images to ensure they remain synchronized (accurate) for a future point of failure. See column 11, lines 38-46. Accordingly, either controller can provide immediate and reliable failover control for the disk storage system.

Nelson et al. is directed to a mirrored memory dual controller disk storage system. Although it discloses maintaining consistency between dual controllers, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Nelson et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

5. U.S. Patent No. 6,260,124 B1

The patent to Crockett et al., US 6260124, discloses a system and method for dynamically resynchronizing backup data. Backup storage 106 is resynchronized to primary storage 104, ensuring that any new updates received during resynchronization are applied in the proper order relative to resynchronization data. Under normal operations, a data mover 114 mirrors data stored in primary storage 104 to backup storage 106. If an error condition arises, preventing mirroring, the data mover 114 stores newly received data in primary storage 104 without mirroring the data to backup storage 106. The data mover 114 also identifies this data in an update map 118. When the error condition ends, the data mover 114 performs a static resynchronization process, serving to update the backup storage 106 with the un-mirrored data, identified in the update map. See column 7, lines 9-25. When new data is received during static resynchronization, a dynamic resynchronization process is invoked to accurately process the updates. Dynamic resynchronization ensures that newly received data records are copied to backup storage in the proper order (if at all) with respect to

versions of the same data being processed by static resynchronization. See column 7, lines 26-55. The host may continue to direct new data to the primary storage even during static resynchronization. Dynamic resynchronization ensures that this new data is written to the backup storage in the proper sequence relative to static resynchronization data. See column 4, lines 34-48.

Crockett et al. is directed to dynamically resynchronizing backup data. Although it discloses a way to ensure newly received data records are copied to backup storage in the proper order, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Crockett et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

6. U.S. Patent No. 6,543,001 B2

The patent to LeCrone et al., US 6543001, discloses a method and apparatus for assuring data consistency in a data processing network including local and remote data storage controllers interconnected by independent communication paths. The remote storage controller or controllers 24, 27 normally act as a mirror for the local storage controller or controllers 22, 23. If, for any reason, transfers over one of the independent communication paths 25, 26 is interrupted, transfers over all the independent communication paths to predefined devices in a group are suspended thereby assuring the consistency of the data at the remote storage controller or controllers. When the cause of the interruption has been corrected, the local storage controllers are able to transfer data modified since the suspension occurred to their corresponding remote storage controllers thereby to reestablish synchronism and consistency for the entire dataset. See column 13, line 62 to column 14, line 25.

LeCrone et al. is directed to a technique for maintaining data coherency. It does so by suspending all independent communication paths when one of the paths is interrupted. It does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, LeCrone et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

7. U.S. Patent No. 6,578,120 B1

The patent to Crockett et al., US 6578120, discloses a method for assuring consistency between a primary volume 29 and a remote secondary volume 33 wherein a host CPU 2 manages track-to-track transfers using loosely-coupled, storage control unit mediated data paths 11, 13. The method includes initial volume synchronization in which concurrent updates to primary tracks are serialized by the copy progression at the secondary volume according to address (FIGS. 3 and 4). The method further includes volume resynchronization in the event of extrinsic error, fault, or the like in which bit maps and timestamps are used to determine the status of primary tracks either lost while in flight to the secondary volume or updated during a volume suspension interval (FIGS. 8-10). The bit maps and timestamps preserve the most recent version copy order of the tracks on the secondary volume. See column 6, lines 16-34. In volume resynchronization, to preserve consistency between volumes, the secondary storage control unit (SCU) must process groups of primary tracks modified before or during suspension to enable the timestamps to catch up. See column 11, line 66 to column 12, line 5.

Crockett et al. is directed to synchronization and resynchronization of loosely coupled copy operations between a primary and a remote secondary DASD volume under

concurrent updating. Although it discloses synchronization and resynchronization to preserve consistency between two volumes, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Crockett et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

8. U.S. Patent Publication No. 2003/0126107 A1

The published patent application of Yamagami, US 20030126107, discloses techniques for managing storage based replication. Specific embodiments provide techniques for performing system backup and restoring. A first storage subsystem 100a and a second storage subsystem 100b are connected to each other via a path. The first storage subsystem 100a is connected to a first host 110a, and the second storage subsystem 110b is connected to a second host 110b. The backup procedure provides a first logical volume 105a in the first storage subsystem 100a, and a second logical volume 105b and a third logical volume 106b in the second storage subsystem 100b. The restore procedure comprises mounting the third logical volume 106b to the second host 110b. Reading a file to be restored from the third volume 106b, and writing the file to the second volume is also part of the method. Further, the restore procedure includes resynchronizing the first volume 105a with the second volume 105b. In a specific embodiment, the restore procedure further comprises recovering a database onto the first volume 105a, if a database application is run on the first host 110a. In a select embodiment, resynchronizing the first volume 105a with the second volume 105b also includes determining data on the second volume 105b pending copy to the third volume 106b. The data pending copy can be tracked by a pending data bitmap, or other data structure, or tracking means. The method can also include, in some embodiments, marking

write data arriving after the command in a pending data bitmap, thereby tracking which data has been modified. See FIG. 6 and paragraphs [0066]-[0077].

Yamagami is directed to backup and restoring systems that employ a remote mirror and a local mirror (see FIG. 1). It does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Yamagami fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

9. U.S. Patent Publication No. 2003/0135650 A1

The published patent application of Kano et al., US 20030135650, discloses a backup method for a system including a network attached storage 100 including a primary volume 120 and a secondary volume 130, and a backup server 200 connected to the network. The network attached storage 100 performs a resynchronization process when a backup request is received at a time point that consistency of a file system can be guaranteed, the resynchronization process making the contents of the primary and secondary volumes coincide with each other, splits the secondary volume 130 from the primary volume 120 and transfers data of the secondary volume 130 to the backup server 200 while an online operation with the primary volume 120 continues. The backup server 200 stores the transferred data in a recording medium 400. The network attached storage 100 performs again the resynchronization process for making the contents of the primary and secondary volumes coincide with each other, after backup completion. Each time data is written in the primary volume 120, the mirror control module 144 writes the same data in the secondary volume 130. Data write (update difference) in the primary volume 120 is recorded and the mirror control module 144 periodically reflects the recorded update difference upon the

secondary volume 130. An operation of making the contents of the primary and secondary volumes coincide with each other is called resynchronization. See paragraphs [0041]-[0043].

Kano et al. is directed to a backup method involving splitting the secondary volume from the primary volume, and resynchronization after backup completion. It does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Kano et al. fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

10. U.S. Patent Publication No. 2003/0177321 A1

The published patent application of Watanabe, US 20030177321, discloses a method and apparatus for enhancing the performance of storage systems to allow recovery after all types of suspensions in remote copy operations. Data is synchronized after an interruption in transfer between a first storage volume 20 of a primary storage system 12 and a first storage volume 21 of a secondary storage system 14 which also includes a second storage volume 22. After the interruption is detected, at the primary storage system 12, a record is provided of the data written onto the first storage volume 20 of the primary storage system 12, and at the secondary storage volume a record is provided of the data written onto the first storage volume 21 of the secondary storage system 14. Then, at least a partial copy of the record of the data written onto the first storage volume 20 of the primary storage system 12 is written onto the second storage volume 22. Using the copy, the first storage volume 21 of the secondary storage system 14 is synchronized with the second storage volume 22 of the secondary storage system 14. Preferably, each of the primary and secondary storage systems include extra volumes, and at the time of suspension of operation

of the mirrored remote pair, bitmaps will be stored from each member of the pair to one of the additional storage volumes. See paragraph [0044] and FIG. 4. These bitmaps may then be used to resynchronize a new pair, even if the information contained on one member of the former pair has been lost. The resynchronization is accomplished by exchanging bitmaps between the new pair, and on that basis determining a set of write operations necessary to resynchronize the new pair. These write operations are then carried out, resulting in a new synchronized pair. See paragraph [0048].

Watanabe is directed to synchronization of multiple remote storage after remote copy suspension. Although it discloses synchronization after suspension, it does not disclose maintaining updates to the data at each of the first and second locations while the link is down, and synchronizing the locations upon reestablishment of a communications link. More specifically, Watanabe fails to teach revising the image of the first data to account for the updates; revising the image of the second data to account for the updates; and either copying data between the primary volume and the secondary volume when the link is restored to thereby resynchronize the primary volume and the secondary volume (as recited in independent claim 1) or revising the data on the secondary storage volume to match the data on the primary storage volume after the link is restored (as recited in independent claim 7); or a sub-system that stores updates to the first data after the interruption in the communications link; and revises the image of the first data to account for the updates to thereby provide a complete first data image for use when the link is restored (as recited in independent claim 12).

(f) In view of this petition, the Examiner is respectfully requested to issue a first Office Action at an early date.

Respectfully submitted,



Chun-Pok Leung  
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP  
Two Embarcadero Center, 8<sup>th</sup> Floor  
San Francisco, California 94111-3834  
Tel: 650-326-2400; Fax: 415-576-0300  
Attachments; RL:rl

60600517 v1